

BACKPLANE SYSTEM USING INCIDENT WAVEFORM SWITCHING

Background Of The Invention

5 1. Field of the Invention.

The present invention relates to system interfaces for backplane technology. More particularly, the present invention relates to bridges and switching systems for establishing transmission interfaces among circuit boards or cards installed in computing device backplanes. The system enables the
10 insertion of a greater number of cards in a backplane topology than has heretofore been possible.

2. Description of the Prior Art.

Standards have been established for the architecture of the hardware employed to enable the exchange of electrical signals among processing devices.
15 The processing devices include integrated circuit systems built on and using printed circuit boards by an increasingly wide array of suppliers. The architecture standards ensure that the various devices will, in fact, be able to communicate with one another as well as with central processing units that control the operation of such peripheral devices. These peripherals include, but are not limited to,
20 printer interfaces, video, audio, and graphics interfaces, memory, external communications interfaces, or any other sort of discrete device performing particular computer-related functions.

The circuit boards associated with the peripherals may be activated upon connection with a primary hardware board, often referred to as a motherboard.
25 The motherboard establishes the interconnection of the central processing unit, power, memory structures, and a backplane bus. The backplane bus is a primary communication interface coupling line having connections to one or more slots or sockets in parallel into which the peripheral circuit boards may be inserted. The slots include physical connectors and input/output interfaces to establish
30 reception and transmission of signals among all devices coupled to the motherboard through the backplane bus. It is the architecture of the backplane

bus that establishes the interface architectures required for the peripheral boards so that communication can occur between all peripherals and the central processing unit in an organized manner.

Several communication hardware protocols have been developed, a number of which have been or are being phased out as being inadequate to support the faster signal exchange rates and increased bandwidth required by newer applications. One of the first such architecture protocols that remains in use on older motherboards is the Industry Standard Architecture (ISA). ISA is an expansion bus slot configuration that accepts plug-ins for peripherals including sound and video displays, for example. Earlier ISA slots were of 8-bit configuration but they are mainly 16-bit slots now. A modified version of ISA, Extended ISA (EISA) was developed to extend the bus capacity to 32-bit with essentially the same convention applied to the ISA backplane slot architecture. Unfortunately, it is designed to run at the relatively slow ISA rate of 8MHz clocking. In today's computing world, that is often too slow.

Although ISA had been the primary standard, the Peripheral Component Interconnect (PCI) bus is now the most frequently implemented interface architecture. Although it continues to replace ISA, that architecture will be used by peripherals designers for as long as computing devices including ISA slots remain in use. PCI provides a relatively high-speed data path (33 MHz) among connected boards. PCI Local Bus is an open architecture specification that allows applications designers greater freedom in interface formation. PCI provides "plug and play" capability in that any peripherals with PCI-based interfaces are automatically configured at start up and therefore generate little to no delay for the central processing unit to establish communication. One advantage of PCI is that it permits the sharing of IRQs. An IRQ (Interrupt ReQuest) signals the central processing unit that activity associated with a peripheral device has started or ended. Since motherboards are configured with only a limited number of IRQ lines, any interface architecture that requires dedicated IRQ lines is necessarily limited. Such is the case with ISA, but not with PCI. For that reason, PCI architecture is ordinarily the first choice.

While the PCI design is an improvement on ISA, it has a history of problems with edge connectors, poor thermal characteristics, and limited interface (input/output – I/O) capability. In fact, PCI backplane configuration is limited to only four slots. That is due in part to the need to resolve signal reflections or noise at the board-to-backplane interface. Subsequent developments to double the clock rate to 66MHZ and to increase the bus width accommodated some of the reflection problems, resulting in a doubling to eight of the number of available PCI slots. That improvement established a new slot architecture referred to as eXtended PCI (PCI-X) architecture. Unfortunately, while providing more slot availability, PCI-X continues to suffer from the same problems associated with PCI.

Sub (B1) Because of interconnection limitations associated with the PCI architectures, there are limits on the number of peripheral devices that may be coupled together. Bridging together two PCI-based backplanes may increase slot capacity; however, that simply increases the size/number of the computer device required to establish desired functionality. It also introduces its own latency and transmission complications. As a result, the Personal Computer Industrial Computer Manufacturing Group (PICMG) developed a standard to address these problems with the PCI functionality. The PICMG combined the architecture of the Eurocard interface with a passive backplane (that is, no active devices to regulate signal propagation, only passive elements), and a high-quality, high-density pin-and-socket arrangement to make improvements. All motherboard components are hence moved from the now passive backplane to a Single Board Computer (SBC) card to be present in the system slot of the passive backplane. The relatively new connection/architecture, identified as CompactPCI, a registered trademark of PICMG (hereinafter referred to as cPCI) improved the peripheral board-to-backplane impedance match, thereby reducing unwanted reflections at that interface. cPCI is designed to be a more robust interface connector to establish solid electrical connections.

The improvements established in cPCI generated the ability to provide eight PCI slots at a 33MHz-clock rate. However, at higher clock rates the number

of available slots is reduced because of the reflective wave signaling technology being used. Using reflected wave technology means the signal travels at half its intended amplitude until it reaches the end of the backplane where it doubles to its intended amplitude and propagates back down the backplane to its point of origin. This round trip delay physically limits the backplane to the number of slot connections it can have for a given clock frequency. The cPCI architecture at 66MHz available in typical computing systems has thus been limited to five open slots.

What is needed is a backplane architecture and interface system that allows for a greater number of slots available for peripheral connections without increasing the footprint of the backplane. Further, what is needed is a backplane architecture and related system that expands the number of slots available without requiring bridging from one backplane to another. Yet further, what is needed is such a backplane architecture and related system that does not compromise the integrity and rate of signal reception and transmission. Still further, what is needed is such a backplane architecture and related system that may optionally be compatible with legacy interface architectures including, but not limited to, ISA.

Summary Of The Invention

It is an object of the present invention to provide a backplane architecture and interface system that provides for an increase in the number of slots available for peripheral connections without increasing the footprint of the backplane. It is also an object of the present invention to provide a backplane architecture and related system that expands the number of slots available without requiring bridging from one backplane to another. It is another object of the present invention to provide a backplane architecture and related system that does not compromise the integrity and rate of signal reception and transmission. Still further, it is an object of the present invention to provide a backplane architecture and related system that may optionally be compatible with legacy interface architectures including, but not limited to, ISA.

These and other objects are achieved in the present invention by establishing a backplane bus driver arrangement that produces incident wave switching rather than reflective wave switching. That is, the driver establishes a voltage swing or step that is of sufficient amplitude to ensure that all receivers coupled to the bus recognize that swing as a valid change of logic state on the first signal edge. Presently, much of the signal propagation associated with PCI architecture is generated at the board level through CMOS, TTL, and LVTTTL drivers and receivers, and their equivalents. Those drivers/receivers generate voltage swings that exceed 1.5 volts in order to produce changes in logic state. When those swing requirements are combined with existing bus impedances related to the connector traces, and the variations in impedances associated with different load impedances, impedance mismatches result. That is particularly the case as increasingly faster switching rates are desired. Those faster switching rates further exacerbate the impedance mismatches, thereby increasing signal noise. The resultant noise must be accounted for so that there is no confusion on the bus. This is achieved by reducing the wave path defined by the bus traces so that the noise settles in time to provide a voltage on the bus that is sufficient to provide a clear logic signal to all receivers on the bus. The path reduction limits the number of connector slots on the motherboard.

The reflective wave propagation is a limit on the bandwidth of the existing motherboard architecture. Unfortunately, it is a common limitation since existing drivers/receivers requiring such reflections are in common commercial use. The present invention involves the application of driver circuitry that generates incident waves rather than reflective waves. One example of such circuitry is the Gunning Transceiver Logic Plus (GTLP) transceiver.

The function of the incident wave switching driver is to ensure a valid change of state on all receivers connected to a backplane bus on the first pass of a propagated transition. For that, the potential of the signal must be strong enough to generate a logic HIGH (V_{IH}) on a rising edge and a logic LOW (V_{IL}) on a falling edge as each of the connected receivers dictates. If the signal must reflect

from the end of the backplane traces to reach a valid potential for signal transition, the driver does not produce incident wave switching. Without incident wave switching, the connector slots must be sufficiently close to the end of the bus to ensure that the reflections generate enough switching potential within a required propagation rate. Therefore, incident wave switching enables longer bus connections without sacrificing signal propagation rate.

The driving circuitry used to generate incident wave switching must be selected with the dynamic characteristics of the backplane in mind. Specifically, variations in capacitance, inductance, and resistance of the bus and the receiver connections affect whether that form of switching is achieved. That is, the driving circuitry must either overcome or balance that variability so that impedance mismatching does not occur, or that it does not generate signal reflection. That resolution of dynamic variability is achieved in the present invention through the selection of appropriate signal driving circuitry and the coupling of appropriate impedance devices between the driving circuitry, its power supply, the bus, and, if required, the individual receivers associated with the individual cards.

As indicated, the driver used in the present invention must be capable of generating incident wave switching. One type of circuit for doing so is described in pending U.S. Patent Application No. 09/132,595 of Oscar Freitas, filed August 11, 1998, and assigned to a common assignee. The content of that application is incorporated herein by reference. One example of such a driver is the GTLP transceiver, such as model GTLP18T612 offered by Fairchild Semiconductor of South Portland, Maine. This driver (transceiver) slows the signal transition enough to minimize noise significantly and therefore avoids reflections.

In addition to the introduction of an incident wave-switching driver, the present invention includes a fairly specific backplane layout to assist in minimizing impedance mismatches. Specifically, the stub impedance (impedance of the connectors for connecting a daughter card to the backplane bus), and backplane impedance. In general, the present invention shortens the connector stubs (thereby reducing their impedances), while maintaining the backplane impedance at or near its current typical value. The present invention relates to selecting

appropriate values for those characteristics of the layout and coupling that with suitable incident wave driver technology to increase the number of slots available on a cPCI backplane or PCI/PCI-X motherboard layout without compromising signal integrity or throughput. It further includes resolution of appropriate signal clocking to minimize clock skew.

As earlier noted, most of the existing peripheral devices employ driver circuitry that generates reflective wave switching. The fabrication of such devices on cards containing such drivers is well defined and well established. For that reason, it would be difficult to conveniently convert those processes to substitute incident wave drivers that would be compatible with the backplane layout improvements described. The present invention includes an interface bridge or interposer that addresses that difficulty. Specifically, the present invention includes an interface circuit that translates the driver/receiver switching architecture associated with each daughter card into driver/receiver switching architecture that produces incident wave switching. For example, the interface circuit may be a TTL-to-GTLP converter that produces the signal swing required to present a valid V_{IH} on a rising edge and a valid V_{IL} on a falling edge during a first propagation on the backplane bus.

The present invention includes a new backplane layout. That combined with incident wave switching and an interface circuit to enable incident wave switching produce a decided advantage in the number of cPCI slots that may be supplied on a backplane without increasing the board footprint and without bridging between or among backplanes. These and other advantages of the present invention will become apparent upon review of the following detailed description, the accompanying drawings, and the appended claims.

Brief Description of the Drawings

FIG. 1 is a perspective view of the backplane system of the present invention, showing 21 slots on a 19" IEEE 1100.10 standard footprint.

FIG. 2 is a simplified schematic diagram of the backplane system of the

present invention, showing the incident wave switching interfaces and the related stub impedances and backplane impedances suitable for minimizing impedance mismatches.

5 FIG. 3 is a layout of the preferred clock path timing for the incident wave switching interfaces suitable for minimizing clock skew.

10 FIG. 4 is a simplified representation of the coupling of a cPCI backplane and a peripheral card that is cPCI configured through an interposer card of the present invention.

15 FIG. 5 is a simplified diagram of the coupling of a cPCI backplane and a peripheral card that is configured through a standard cPCI bridge coupled with a state machine for conversion of a reflective wave driver system to an incident wave driver system.

20 FIG. 6 is a simplified diagram of a coupling of a cPCI backplane and a peripheral card that is configured through a modified cPCI bridge with state machine for coupling to an incident wave driver system.

Description Of The Preferred Embodiments Of The Invention

25 A backplane system **10** of the present invention is shown in FIG. 1. The system **10** includes a backplane bus **11** having matched impedance terminations **12** and **13** on the back or south side of the bus **11**. The impedance values of the terminations **12** and **13**, the impedance values of the slot connectors identified herein, and the impedance values of the backplane bus traces **11**, together provide the backplane impedance of the system **10**. The impedance terminations **12** and **13** are coupled to a supply rail that establishes the maximum swing potential of the backplane bus **11**. The system further includes a system clock
30 card **14** that generates the clocking scheme for each of one or more of a plurality of peripheral devices provided on daughter cards such as near card **15** and far

card **16**. Each card includes one or more integrated circuits for performing functions of interest regulated by one or more central processing units tied to the backplane bus **11**. The cards **15** and **16** are of the type that may be inserted into conventional cPCI backplane slots such as slots **17a-17u** shown. The slots are
5 located on the front or north side of the bus **11**.

The system **10** further includes one or more interposer cards such as interposer cards **18** and **19** that translate reflective wave switching associated with the conventional peripheral devices such as cards **15** and **16**, into incident wave switching. The interposer cards **18** and **19** include translation circuitry to
10 convert TTL-based potential swings into narrower potential swings, such as those associated with GTLP transceivers, for example. The interposer cards **18** and **19** further include logic and control circuitry for regulating the propagation of converted signals to and from the backplane bus **11**.

As illustrated in FIG. 2, the system **10** includes for each slot an interposer
15 card, such as interposer cards **18** and **19**. Each card in a slot of the backplane bus **11** includes a transceiver **20** that receives or transmits input/output signals from or to the bus **11**. A suitable transceiver **20** is the GTLP18T612 driver offered by Fairchild Semiconductor of South Portland, Maine. That transceiver converts all bus signals to GTLP levels, along with any control signals. System card **14**
20 includes an incident wave compatible clocking scheme, such as a GTLP clocking device **21** for use with GTLP transceivers **20**. Preferably, the GTLP clocking device is the GTLP16C816 clock driver offered by Fairchild Semiconductor Corporation of Portland, Maine. That clocking device **21** drives and receives the GTLP clock and it generates the clocks for all of the peripheral cards connected
25 through an interposer to the bus **11**, such as cards **15** and **16** of FIG. 1.

FIG. 3 represents a preferred clock path-timing scheme to minimize clock skewing. As can be seen, sets of backplane slots of the bus **11** are coupled through their respective transceivers to the system card **14** that includes the GTLP clock driver. Each pathway for the respective sets of slot drivers is
30 preferably of the same length. In that way, clock skew across the equivalent clock sets is reduced. The number of slots tied to each clock driver of the system

card **14** is dependent upon pin spacing size and spacing limitations and the particular driver employed. For the GTLP16C816 clock driver, the number of slots coupled in parallel with a particular clock driver is preferably three.

Returning to FIG. 2, the configuration of the backplane bus **11** and the impedance elements connected thereto is important in minimizing mismatches that would prevent incident wave switching. It has been determined through evaluation that in order to achieve that goal, the terminals **12** and **13** should be about 0.5 inches in length with an impedance of about 65 ohms, and the connector or stub length for the individual slots should be 0.75 inches, with an impedance of about 50 ohms. When loaded, the impedance of the bus **11** may be between about 20 ohms and about 30 ohms but for optimal operation the system impedance provided by terminations **12** and **13** is about 40 ohms.

The interposer card is preferably implemented as an interposer state machine and the GTLP transceivers. As illustrated in FIG. 4 for a cPCI backplane architecture and a cPCI configured peripheral card **22**, the interposer is coupled directly between the peripheral card **22** by a card connector **23**. The interposer card is similarly coupled to the slot of the backplane **11** by way of an interposer connector **24**.

As shown in FIG. 5, with a peripheral card **25** having cPCI architecture, a cPCI interface **26**, such as any of those commonly programmed into Altera's APEX family of devices, is used to establish an otherwise conventional interface between the card **25** and the backplane **11**. The interface **26** is modified such that it includes interposer state machine circuitry **27** for timing, addressing and arbitration, as an ASIC, a complex programmable logic device (CPLD), or as a field-programmable gate array (FPGA) configured to control operation of the transceivers **28**. The state machine **27** may be coupled as a separate integrated circuit to the interface **26** or it may be fabricated on the same chip as the interface **26** as shown. That is, the state machine **27** may be incorporated into the fabrication of the interface **26**. Optionally, it may be fabricated as part of the fabrication of the interposer card. It may also be a stand-alone device connected to the relevant circuitry of the interface **26**. It is to be noted that the transceivers

28 may also be stand-alone devices physically connected to the relevant circuitry of the interposer card either on the same circuit board or on another board.. They may also be fabricated in silicon integrally with the fabrication of the associated circuitry of the interposer card or the interface.

5 In an alternative embodiment of the interface shown in FIG. 6, the peripheral card **25** having cPCI architecture, the cPCI interface **26** of FIG. 5 is fabricated as an integrated modified interface **29** including interposer state machine circuitry **30** and the transceiver **31** forming part of a comprehensive integrated incident wave interface **32** used to establish an otherwise conventional interface between the card **25** and the backplane **11**. The interface **29** is modified with the state machine circuitry **30** for timing, addressing and arbitration, as an ASIC, a CPLD, or a FPGA configured to control operation of the integrated transceivers such as integrated transceiver **33**, one for each slot of the backplane bus **11**.

15 Using representative operating specifications for the GTLP18T612 transceiver listed in Table 1 and applying them to the cPCI specifications listed in Table 2, it can be seen that the present invention enables design of a 21-slot backplane at 33 MHz and an alternative 14-slot backplane at 66 MHz. The calculations for that analysis show that the backplane delays associated with the incident wave switching permit substantially more slots.

Parameter	Value	Units
$T_{\text{Prop delay TTL to GTLP out}}$	6.3 max	Nanoseconds
$T_{\text{Prop delay GTLP to TTL out}}$	5.8 max	Nanoseconds
$T_{\text{Setup TTL to clock}}$	1.1 max	Nanoseconds
$T_{\text{Setup GTLP to clock}}$	3.0 max	Nanoseconds
$T_{\text{Prop delay clock to GTLP out}}$	6.5 max	Nanoseconds
$T_{\text{Prop delay clock to TTL out}}$	4.6 max	Nanoseconds

Table 1 Summary of GTLP18T612 Specifications

Parameter	PCI-X @133	PCI-X @100	PCI @66	PCI @33	Unit MHz
T_{val} (max)	3.8	3.8	6.0	11.0	Nsec.
T_{prop} (max)	2.0	4.5	5.0	10.0	Nsec.
T_{skew} (max)	0.5	0.5	1.0	2.0	Nsec.
T_{su} (min)	1.2	1.7	3.0	7.0	Nsec.
T_{cyc}	7.5	10.0	15.0	30.0	Nsec.

Table 2 Summary of PCI Timing Budget

5

In the first scenario for 33 MHz backplane operation, assume a PCI-X compliant peripheral card with cPCI hardware designed for 100 MHz operation (resulting in the equivalent of the timing budget for 100 MHz operation in Table 2 above) coupled to an interposer card of the present invention that includes the GTLP18T612 transceiver and an appropriate clock driver. The values presented in Table 1 apply to a slot pitch of 0.8 inches. Therefore, for a 19" IEEE 1100.10 backplane architecture, that is the equivalent of a 21-slot backplane. The total time delay value associated with worse case TTL-to-GTLP output for that backplane arrangement equals 3.8 + 6.3 nanoseconds. That is 0.9 nanoseconds more than the budgeted T_{val} for cPCI @ 33 MHz. The combined setup time is 1.7 + 5.8 nanoseconds. That is 7.7 nanoseconds or 0.7 nanoseconds over T_{su} for cPCI @ 33MHz. The timing budget is therefore 10.0 - 0.9 + 0.7 nanoseconds = 9.8 nanoseconds. Since system propagation delay is 10 nanoseconds for cPCI @ 33MHz, the calculated propagation delay available is within the system delay budget and a complete 21-slot backplane can be implemented.

In the second scenario again for 33 MHz backplane operation, assume a cPCI compliant peripheral designed for 66 MHz operation (resulting in the equivalent of the timing budget for 66 MHz operation in Table 2 above) coupled to an interposer card of the present invention that includes the GTLP18T612

transceiver and an appropriate clock driver. Again, the values presented in Table 1 apply to a slot pitch of 0.8 inches and so the maximum backplane architecture for a 19" IEEE 1100.10 backplane would be the equivalent of a 21-slot backplane. In this scenario, the total time delay value associated with worse case TTL-to-GTLP output for that backplane arrangement equals 6.0 + 6.3 nanoseconds or 12.3 nanoseconds. That is 1.3 nanoseconds over the budgeted T_{val} for cPCI @ 33 MHz. The combined setup time is 3.0 + 5.8 nanoseconds. That is 8.8 nanoseconds or 1.8 nanoseconds over T_{su} for cPCI @ 33MHz. The timing budget is therefore 10.0 + 1.3 + 1.8 nanoseconds = 13.1 nanoseconds. Since system propagation delay is 10 nanoseconds for cPCI @ 33MHz, the calculated propagation delay is a portion of the system delay budget for a 21-slot backplane. In particular, the available budget is about 69% of that associated with a 21-slot backplane. Therefore, a backplane having 14-slots (21x69%) can be implemented for the indicated peripheral card I/O hardware.

It is to be noted that the two scenarios described above make a 21-slot and a 14-slot backplane architecture available without modification to the clocking associated with the selected transceiver used to produce incident wave switching.

However, by conventional timing regulation means, the transceivers associated with the slots may be re-timed so that the interposer card may be employed to generate a cPCI solution at 66 MHz for a 14-slot backplane. Specifically, latching the signal transmissions at the transceiver regulates signal propagation. By latching the transceivers on the rising edge of the clock rate, the signal transfers are delayed by two clock edges (one to acquire into the transmitting GTLP transceiver and one to latch into the receiving GTLP transceiver). The delays are pipelined onto the backplane so that only two latency cycles are required at the start of a burst transfer. In order to do this, the state machine previously described may have to be modified in a manner known by those skilled in the art so that the delays can be accommodated. Through the noted retiming, a 66 MHz backplane may be implemented with far more slots available than is currently available under existing cPCI and PCI-X specifications.

While the invention has been described with reference to particular example embodiments, it is intended to cover all modifications and equivalents within the scope of the following claims.